

Eye on Innovation

Manufacturing Edition: SGI News

Two Moments of Truth

“We use the SGI supercomputer to help us develop virtual prototypes that now replace the ‘initial’ physical prototypes of both products and production systems. These virtual prototypes provide a timely and cost-effective means by which P&G can determine the ‘fit, work, and financial sense’ components used to evaluate a product—before committing to building a physical representation.”

— Tom Lange, Associate Director for Modeling Simulation and Analysis, Procter & Gamble.

Procter & Gamble uses CAE and SGI® supercomputers to ensure product and production system innovations fit, work, and make financial sense

Procter & Gamble (P&G) has been a manufacturing stalwart for 165 years and counting. The package goods giant has annual revenue of \$40 billion and spends close to \$1.6 billion each year on research and development—the highest in the consumer package goods industry.

P&G products include nearly 300 of the world’s most recognized brands, from Ivory® Soap, Pringles®, Charmin®, and Tide® to Downy®, Crest®, Mr. Clean®, Pampers®, and many more. Annually, P&G products find their way into the homes of 5 billion consumers in 160 countries.

So why does P&G need a supercomputer to design and manufacture things as simple and mundane as toothpaste and toilet paper? The answer is complex.

A Technology Company That Thrives on Innovation

Despite its identity as a consumer package goods company, P&G, at its core, is a technology company, one that thrives on innovation. P&G holds over 28,000 patents and is granted a new patent at a rate of more than one a day.

Tom Lange, a 25-year P&G veteran, is responsible for enlisting computer-aided engineering [CAE] technologies that ensure the company’s innovations see the light of day. These days, Lange wears multiple hats. Officially, he is the company’s associate director for modeling simulation and analysis for corporate engineering. His other titles within P&G include chief technologist for reliability engineering and head of computer-aided engineering.

“Innovation is the lifeblood of our business,” says Lange. “When we create new products our mantra is simple—it has to fit, on the body, in the hand, or on

the plant floor. It has to work, i.e., do what it’s supposed to do. Finally, it has to make financial sense—we have to make it faster, more cost-effectively, and it must add value to the company.”


Helping Lange and his design, engineering, and research teams make those innovations happen is a new SGI® Origin® 3800 system with 128 CPUs and 128GB of memory. The SGI supercomputer allows P&G to create and test prototypes of products and the machines that manufacture them in a virtual state—thereby eliminating the need for first-round physical prototypes. The supercomputer also allows P&G teams to study the effects of production-line changes without building the actual equipment. This eliminates costly and time-consuming trial-and-error testing that was once done using physical prototypes.

The Origin 3800 system is used by P&G analysts worldwide to assist in modeling and simulation projects.

The Challenge

CAE first appeared as a tool for the defense industry before migrating to aerospace endeavors [military and commercial] and eventually durable consumer goods [automobiles, farm machinery]. Says Lange, “The use

[continued on page 3]



IN THIS EDITION

1 Customer
Procter & Gamble—
Creating Better Products
for Consumers

5 Technology
MCAE Application
Productivity with
SGI® Altix™ 3000

7 Partner
CEI—Bringing into
View What Was Once
Out of Sight

10 Industry
SGI Gives Defense
Industry the Power
to See Beyond

Achieving Customer Satisfaction and Measurable ROI through Design for Six Sigma

Faced with growing globally competitive pressures and more-demanding customers, manufacturing companies are increasingly embracing quality-boosting efforts for improving customer satisfaction and reducing costs. Spurring the companies towards this move is the realization that only those with the finest quality will thrive in the new competitive environment. Quality alone is not the sole driver, but lower costs and improved productivity are almost always by-products of making things better.

Six Sigma, in the context of product design or service-oriented business, is a strategy consisting of analytical-based practices and procedures aimed at reducing the number of product defects and variability resulting from any process to achieve a near-perfect quality level [specifically, the value of Six Sigma represents 3.4 defects per million opportunities]. GE and several other major corporations have instituted company-wide programs to make Six Sigma the company culture and the way they work.

The traditional Six Sigma approach, while being very efficient, is applied in the later stages of a product life cycle, i.e., manufacturing and production, to identify and fix defects and is hence a reactive quality strategy. While defects may be easy to identify in the later stages of a product life cycle, they can be costly to correct. Hence, with the traditional Six Sigma approaches, companies typically run into a “4 or 4.5 sigma barrier” and must resort to a proactive design strategy to avoid major defects and quality issues.

Realizing that a majority of the downstream producibility and quality problems are a result of the earlier design phase activities and get locked in early during the product development cycle, the focus of DFSS [Design for Six Sigma] is on the design and development stages of the product life cycle. DFSS allows engineers to predict production and performance capability early in the product development process and improve it, rather than reacting to poor production capability and performance after the fact.

The goal of DFSS, therefore, is to create new game-changing products and services that have quality designed in from the start—the way for companies to realize the full benefit of Six Sigma performance.

DFSS has a significant effect on longer-term profitability through improved products, which results in increased customer satisfaction, hence improved market share and increased profits.

The GE Medical Systems LightSpeed™ CT Scanner project in 1998 was GE’s first DFSS system, made full use of Six Sigma/DFSS tools, and is considered the “biggest breakthrough in CT in a decade,” according to Dr. Gary Glazer of Stanford. A disciplined systems-engineering approach resulted in meeting 90 critical-to-quality and critical-to-customer requirements, such as image quality, speed, software reliability, and patient comfort. The CT scanner involved several leading-edge technologies including the world’s first 16-row CT detector, multislice data acquisition, 64-bit RISC computer architecture, and the long-life Performix tube. The DFSS process resulted in better image quality and much faster scanning as well as a two-year reduction in development time and higher market share.

For GE, Six Sigma has had a major impact on the bottom line. In 2001, the margin on revenue growth from customer satisfaction, including DFSS products, is estimated to be \$1.5 billion.

DFSS of complex systems, such as in the automotive, aerospace, and general manufacturing industries, results in a simulation environment with the following characteristics: a high number of variables, a substantial number of design subsystems, interdependency and interaction between the subsystems, and large, complex models across several engineering disciplines. These attributes are representative of a systems engineering environment that would benefit from the use of high-performance computing, complex data management capabilities, and high-end visualization.

Combining near-perfect manufacturing quality and a real focus on the customer is a tremendous advantage when introducing new products. SGI products can really make a difference by allowing the optimization of product design, coupling product design to manufacturing capability, working with the customer on product visualization, and a whole host of additional benefits. This in turn can allow SGI’s customers to produce products with superb quality and great customer satisfaction.

Commentary printed by SGI with permission.



By Dr. Lewis S. Edelheit, senior vice president of corporate research and development for General Electric Company [retired] and member of SGI’s board of directors

Eye on Innovation

CUSTOMER

[continued from page 1]

of CAE is almost ubiquitous in those industries. My job is to bring it to P&G's consumer package goods and the machines used to make them."

The challenge that Lange and P&G face in using CAE to develop consumer package goods is multifaceted. According to Lange, "Consumers see the package goods created by P&G as relatively inexpensive per-use products. What they don't see is the complex technology required to develop and keep the products inexpensive."

There is an inverse relationship that exists in consumer package goods between the consumer per-use cost of a product and the complexity of the technology required to produce it. "A perfect example," Lange points out, "is Charmin bathroom tissue. Charmin has very important product attributes for which consumers are willing to pay, but only within reason. The per-use cost has to be very low for it to be an attractive purchase. However, the machines used to manufacture Charmin cost hundreds of millions of dollars apiece. And in terms of complexity, the number of lines of computer coding required to run each machine and the number of stationary and moving parts required to build and operate the machine make it every bit as complex as some of today's more advanced commercial aircraft."

Considering the number of brands P&G markets, it's easy to see how the cost of introducing innovation to both new and long-standing brands can multiply exponentially. Lange adds, "CAE has been a boon to P&G innovation, but because every product we manufacture is unique unto itself, there is no single recipe for applying CAE to product design and development."

Two Moments of Truth

There are, however, two moments of truth that drive P&G innovation. The first is when the consumer decides to buy a P&G product. If it isn't priced right or has no perceived value for customers, they won't try it. Hence, the costs of materials that go into a product as well as the cost to manufacture the product are paramount in the first moment of truth. CAE has a strong role in evaluating suitable materials for P&G products, to ensure that they work, and in developing manufacturing systems that help keep them affordable. CAE helps to ensure that product containers don't break or crack when dropped, viscous fluids flow easily from their containers, and lids of every type don't leak.

The second moment of truth occurs when consumers open the package and use the product. It has to perform as intended and it has to live up to its promise, whether it's laundry detergent, disposable diapers, or a

snack food container. Here, again, CAE plays a role in determining how P&G products function during actual consumer use.

Explore Digitally, Confirm Physically

To develop package goods in a physical-only environment is not only costly, it's also an innovation killer. No company, package goods or otherwise, is interested in introducing something dramatically different to consumers without significant testing. Before CAE, bringing a new product to market was a multiyear exercise in trial and error.

To avoid the cost and time constraints of creating products in the physical world, P&G chooses increasingly to "explore digitally, confirm physically," according to Lange. "We use the SGI supercomputer to help us develop virtual prototypes that now replace the 'initial' physical prototypes of both products and production systems. These virtual prototypes provide a timely and cost-effective means by which P&G can determine the 'fit, work, and financial sense' components used to evaluate a product—before committing to building a physical representation. Virtual prototyping gives us a chance to ask what-if, and then test it in any number of ways to determine next steps, if any, moving forward."

Despite the obvious benefits of virtual prototyping, Lange is quick to add, "P&G will continue to make physical prototypes in the latter stages of product development. They still carry many intrinsic tactical consumer findings that simply cannot be replaced virtually."

Less Risk, More Creativity

Consumers have a faster appetite for change than the manufacturers can satisfy cost-effectively. In the case of P&G, CAE and SGI supercomputers are helping to narrow the gap by enabling the company to pursue a wider range of creative solutions to meet consumer needs without having to invest in costly infrastructures. Lange considers the reduction in costs and the increase in innovative opportunities as the major benefits of using CAE and supercomputing to perform modeling and simulation explorations. "We have documented cases where it has saved P&G millions of dollars and months to years in development time, for both products and the machines that make them."

CAE Tools

Many of the CAE tools used by P&G are readily available commercial applications, including LS-DYNA® from Livermore Software Technology Corporation [www.lstc.com], for structural analysis

[continued on page 4]



Eye on Innovation

CUSTOMER

[continued from page 3]

and impact simulation; ABAQUS® from ABAQUS [www.abaqus.com], for stress, heat transfer, and other types of analysis in mechanical, structural, and related engineering applications; and FLUENT® from Fluent, Inc. [www.fluent.com], for evaluating the computational fluid dynamics of product designs.

The company is also using applications derived from national laboratory codes. The reason for this is that most consumer package goods are not made out of metal and, therefore, have a very different set of dynamics than products made of aluminum or steel.

Lange points to products like shampoo, toothpaste, skin lotions, and detergents, which tend to be sticky, goopy, complex, and unusual liquids. “These are all nonlinear materials, and they dominate our environment. We need to know how they perform,” he says. “How do they react to heat and cold? Do they flow easily from their containers? How do they affect the surfaces with which they come in contact?”

Consumer package goods also have complex geometries. They are ergonomic by design, and how people interact with them is important to P&G’s second moment of truth—does the product perform as intended in the hands of the customer? Modeling and simulation tools are also helping P&G evaluate how the range of human variability affects a product’s performance.

Because the consumer package goods industry was not an early adopter of commercial CAE applications, as were defense, heavy industry, and durable goods manufacturers, the available commercial applications tend not to favor this industry. Lange says, “I believe the main reason for this is that the business value of CAE tools to package goods manufacturers is harder to establish, initially. It’s easy for software developers to understand that a simple change in automotive design and development costs can save a manufacturer thousands of dollars per vehicle. But when you talk

about developing ways to make a 24-cent plastic bottle cap more affordable to manufacturers, you get a very different reaction.

They ask you, ‘Why do you care?’

We care because P&G makes a half-billion of them at a time and a one-cent-per-cap savings is significant.”

Moving forward, P&G is looking to work with software developers whose applications will most closely

address the physics and geometric complexities inherent in the design and manufacture of consumer package goods.

Investing in Supercomputing

The cost and time savings realized by P&G’s investment in CAE tools and use of supercomputers to handle complex modeling and simulation is apparent. CAE has moved from research tool status to production requirement tool. The computational workload handled by the company’s SGI supercomputer continues to grow as well. Around-the-clock usage of the supercomputer is standard, with peak utilization sometimes running as high as 80%.

P&G chose to use SGI supercomputers to help advance the company’s products and production systems design for a number of reasons. One of the most significant was versatility. The Origin 3800 system, with 128 CPUs and 128GB of memory, allows Lange to choose between solving several large computational problems and numerous smaller problems or a combination of the two—with a single computing system. The versatility of the Origin 3800 system also allows P&G to support the modeling and simulation needs of its numerous brands located throughout the country.

Another key factor in the decision to go with the Origin 3800 system was performance reliability and the relatively simple IT support structure needed to watch over the system. “I want my analysts to be more worried about physics and meeting our business needs than about how to keep a computer working,” says Lange. The performance reliability that P&G experienced with a Silicon Graphics Onyx2 system prior to acquiring the Origin 3800 system factored into the company’s decision.

Lange says P&G’s selection of the new SGI supercomputer was driven by return on investment.

“Development is a zero-sum game. The money spent on improving R&D infrastructure needs to be equal to the savings realized from no longer having to create initial physical prototypes. We try to make infrastructure investment decisions that will sustain our R&D efforts for at least three years, which means we have to anticipate or predict how much we can replace over that time period.”

“Because the package goods industry deals with unusual materials and complex geometries that require excessive gridding, we require a tremendous amount of problem-solving memory. Large memory and high speed are essential. The more, the better, especially when it comes to speed. In choosing the SGI Origin 3800 sys-

[continued on page 12]

“We have documented cases where it has saved P&G millions of dollars and months to years in development time, for both products and the machines that make them.”

—Tom Lange, Associate Director for Modeling Simulation and Analysis, Procter & Gamble.

Actual

Simulated

manufacturing

MCAE Application Productivity with SGI® Altix™ 3000 Technology

The combined forces of mechanical computer-aided engineering (MCAE) application software and HPC system technology provide engineers with an increasingly competitive advantage in today's global manufacturing market of product development. Manufacturers and suppliers in automotive, aerospace, and a variety of general manufacturing sectors benefit from MCAE applications that enable reduced design-cycle time and costs and overall improvements in design quality.

This article examines the MCAE productivity benefits of a new HPC system technology, the SGI Altix 3000 family of servers and superclusters, which was developed to advance the current capabilities of technical HPC. The design of SGI Altix 3000 combines the open-source 64-bit Linux® operating system, the Intel® Itanium® 2 microprocessor, and the SGI® NUMAflex™ shared-memory system architecture design. With this new server technology introduction, SGI offers the MCAE community a dual-platform roadmap based on UNIX® and Linux.

SGI Altix 3000 Introduction

Recent developments in HPC technology continue to rapidly advance the MCAE simulation capabilities of engineers across all disciplines. SGI is an established and leading supplier of HPC technology for MCAE simulation with the company's SGI® Origin® server family, based on the company's proprietary MIPS® microprocessors and IRIX® operating system. On January 7, 2003, SGI expanded its HPC offering with the introduction of SGI Altix 3000, the industry's first global shared-memory cluster that combines SGI NUMAflex supercomputing architecture with Intel Itanium 2 processors and the 64-bit Linux operating system.

The SGI® NUMA (nonuniform memory access) architecture was introduced in the SGI® Origin® 2000 server in 1995 and later advanced with the SGI NUMAflex modular design concept of the SGI® Origin® 3000 servers. This is the same high-bandwidth and low-latency NUMA architecture that is available in SGI Altix 3000, yet with a significant cost-performance advantage for MCAE applications. This advantage is achieved by, among other factors, contribution to and leverage of investments in Linux by the open-source community, and the Itanium 2 microprocessor roadmap from Intel.

SGI Altix 3000 is recognized by the Linux community as the first Linux cluster that scales to 64 processors and 512GB of shared memory within each single Linux OS-image node and the first cluster of any variety to allow global shared-memory access across nodes. These 64-processor single nodes can be clustered with a choice of scalable interconnect networks, including the proprietary SGI® NUMAlink™ interconnect technology, to much larger system configurations—up to 2,048 processors and a total of 16TB of memory—that make up an SGI Altix 3000 supercluster. The high-bandwidth SGI NUMAlink interconnect fabric of SGI Altix 3000 superclusters delivers information between cluster nodes up to 200 times faster than conventional clustering switches.

SGI Altix 3000 is binary compatible with the industry-standard 64-bit Linux distribution, currently based on the 2.4.19 kernel. In addition, SGI offers differentiated middleware and other functionality to enhance demanding HPC workloads in a bundle called SGI ProPack™. SGI ProPack is a set of user tools that ride on top of Linux and is similar to other commercial software packages. SGI ProPack is used to boost the performance of Linux and user applications on the Altix 3000, not to alter Linux itself. Such enhancements are commonplace in the industry and are also offered by other Linux system providers.

The performance of the SGI Altix 3000 family has set new records with many industry-standard benchmarks. In particular, the SPECfp®_rate_base2000 benchmark, a measure of a systems compute performance, and STREAM Triad benchmark, a measure of memory bandwidth performance, show SGI Altix 3000 easily outperforming all of the top-end UNIX OS-based systems that are popular with MCAE applications today.

MCAE Software and Industry Practice

Rapid progress in MCAE simulation performance has been influenced by advanced developments in both application software algorithms and HPC hardware systems. From a software algorithm and hardware perspective, there are three MCAE disciplines to consider for their requirement on HPC resources: implicit and explicit finite element analyses (FEA) for structural analysis and computational fluid dynamics (CFD) for fluid flow simulation.

[continued on page 6]



[continued from page 5]

The three MCAE disciplines exhibit a range of HPC resource demands and highlight the importance of a balanced HPC system architecture. The features desired most of a balanced system include [1] high-speed processors with large cache; [2] large addressable memory; [3] high memory-to-processor bandwidth rates; [4] high disk-to-memory I/O rates; and [5] a low-latency interconnect that provides efficient parallel scalability to hundreds of processors.

By far the most important HPC advancement in recent years for MCAE applications is the parallel scalability of MCAE software. Most commercial MCAE software employs a distributed-memory parallel [DMP] technique for compatibility across the range of available HPC architectures. Other techniques include shared-memory parallel [SMP] and hybrid parallel schemes that take advantage of both DMP and SMP within a single computation.

While each discipline has inherent complexities with regard to efficient parallel scaling depending upon the particular parallel scheme and HPC system architecture, CFD can scale efficiently to hundreds of CPUs, explicit FEA can scale to more than 50 CPUs, and implicit FEA can scale to up to 10 CPUs.

Single-job turnaround of MCAE simulations on SGI Altix 3000 has been impressive, with many ISV applications showing a performance level that exceeds high-end UNIX systems. For some MCAE applications, Altix 3000 demonstrates efficient parallel scalability as high as 64 CPUs. This stand-alone job performance is critical for success of any new server today, but true commercial success requires the additional capability of job throughput, which more closely captures the industry's true MCAE practice. Industry practice most often combines the use of moderate single-job scalability with multi-job throughput. That is, a single job typically uses 12 CPUs on average, in a mix that combines several disparate jobs that require a throughput level of productivity.

Throughput is the domain and responsibility of hardware vendors, and SGI leads the industry in this critical HPC industry requirement. The NUMAflex system architecture was designed for a combination of single-job turnaround and multi-job throughput. Based on this architecture, SGI Altix 3000 offers MCAE simulations a high-availability, nondegrading, and efficient application environment to ensure that turnaround and throughput are delivered in support of hundreds of simultaneous users with a demanding mix of disciplines.

Multi-job throughput has been demonstrated on SGI Altix 3000 with a range of MCAE commercial software applications and industrial-sized customer jobs that exhibit an average throughput degradation of 10% compared with their turnaround times. This means that a single job, on average, requires just 10% more time to complete when that job is included in the mix of jobs.

SGI Altix 3000 Technology Benefits for MCAE

SGI continues to invest in the MCAE community's shared vision with the introduction of SGI Altix 3000, an HPC server that will enable further advancements for MCAE applications. SGI Altix 3000 delivers the existing application advantages of NUMA with complete 64-bit Linux compatibility, but in a shared memory capability that is not available in conventional clusters.

Unlike conventional clusters, the SGI Altix 3000 scalable memory architecture can conduct memory-resident MCAE analyses with the data as a whole entity, without breaking the data into smaller partitions to be handled by individual processors. Consequently, a programmer or user does not need to spend time developing rules to divide the data into smaller sets, as is required with domain decomposition techniques for distributed parallel MCAE.

The primary reason for efficient multi-job throughput capability with SGI Altix 3000 is the latest Linux kernel scheduler that performs well with large memory, an advantage SGI Altix 3000 has over commodity clusters. The trend of cheap and abundant commodity DRAM means this advantage will grow over time.


Another important advantage is I/O performance, a requirement for throughput of implicit FEA applications that simulate dynamic response of a structure. SGI Altix 3000 offers I/O rates of more than 2GB per second, well beyond the typical Linux barrier of 500MB per second.

MCAE Technology Directions


The design breakthroughs of SGI Altix 3000 offer the industry a shared-memory choice for open-source computing with the cost-effectiveness of conventional clusters. With SGI Altix 3000, the entire suite of MCAE applications can now achieve capability levels that are equivalent to the most mature UNIX platforms, but at less than half the cost. This capability, along with its economic benefits, will further expand the use of MCAE to include a variety of new applications and practices that will become routine for product development.

[continued on page 12]

SPECfp_rate_base2000 benchmark scores [higher is better]



STREAM Triad Benchmark results in GB/sec [higher is better]



Bringing into View What Was Once Out of Sight

SGI and Computational Engineering International (CEI) join forces to provide rapid-fire analyses and interactive visualization that sheds new light on complex engineering data.

Complex engineering models recently have crossed a threshold that less than two years ago seemed impossible to many within the mechanical computer-aided engineering (MCAE) community. The collaborative efforts of CEI, with its tools for the evaluation and review of structural and fluid dynamic simulation data, and SGI, maker of industry-leading high-performance computing (HPC) and advanced visualization technologies, have turned impossibility to deliverability. The sum of their parts—breakthrough speed, detailed analysis, and interactive visualization—now makes 3D models containing hundreds of millions, even billions of cells, a daily reality.

Optimization through Collaboration

CEI, developer of EnSight® and EnSight® Gold postprocessing software for MCAE, and SGI have been working together since 1994 to optimize the performance and scalability of CEI software when used by manufacturers to review finite element analysis (FEA), computational fluid dynamics (CFD), and other common analyses on SGI platforms. Current focus is on the latest SGI® Onyx® family systems with InfinitePerformance™ graphics subsystems.

MCAE Visualization

Manufacturers rely increasingly on MCAE to improve product design and quality. MCAE simulations allow engineers to evaluate the effects of mechanical stress, vibration, impact loading, fluid-induced pressure, and many complex physical phenomena on potential designs while avoiding costly and time-consuming physical prototypes.

A typical MCAE analysis consists of three phases: preprocessing, computation, and postprocessing. Harpoon, CEI's entry into preprocessing, is an automatic hex-dominant mesh generator that can move CAD geometry files to mesh with as little as three mouse clicks, effectively eliminating the workflow bottleneck associated with preparing models for MCAE simulations. Complex CAD models that once took weeks to mesh using traditional methods now take only minutes. After applying load condi-

tions, the computational analysis can be run on a compute server such as an SGI Origin or SGI Altix family supercomputer, followed by the exportation of solver results to EnSight or EnSight Gold for postprocessing on visualization systems such as the SGI Onyx family.

The postprocessing phase is critical to the understanding of an MCAE simulation. According to CEI's president Kent Misegades, "The true value of an analysis is determined both by the accuracy of the analysis code and the ability to interpret all of the significant information contained in the analysis results. Without a strong visualization tool, it might be difficult to recognize the existence of problems or inefficiencies within a given system design, or what is causing a certain physical phenomenon."

EnSight and EnSight Gold are dominant players in MCAE postprocessing for industry-leading companies in aerospace, automotive, defense, materials processing, power generation, structural engineering, and turbomachinery. They also play a significant role in scientific research.

Misegades says about EnSight, "We set out to create a single piece of software that addresses postprocessing needs on an ever-expanding number of levels for a range of manufacturers. Manufacturers are under increasing pressure to reduce the number of required analysis tools. In providing a standardized postprocessing tool set, we've provided an excellent opportunity

[continued on page 8]





Image courtesy of Advantage CFD, www.advantage-cfd.co.uk

"Our experience with OpenGL Vizserver has been impressive, and our software is well matched to the concept of remote visualization. SGI technology and VAN permit easier access to EnSight's powerful visualization feature set.

— Kent Misegades,
President, CEI

Eye on Innovation

PARTNER

[continued from page 7]

for manufacturers to reduce the number of tools. This drives product quality, reduces research costs, speeds product development, and accelerates time to market.”

EnSight

EnSight is the world’s most advanced software tool for the interactive visualization and animation of data from a wide variety of CAE simulations in fields such as CFD, combustion modeling, structural analysis, crash and impact analysis, thermodynamics, electro-magnetics, and others. EnSight enables engineers to read in and work with multiple computational data sets for comparative purposes.

EnSight Gold

EnSight Gold shares EnSight’s broad feature set and takes full advantage of parallel processing [SMP and DMP] and rendering, provides support for an array of virtual reality [VR] environments, and enables real-time collaboration.


EnSight Gold is the direct result of requests from scientists and engineers working on the Accelerated Strategic Computing Initiative [ASCI], a program sponsored by the Department of Energy for simulated nuclear testing being conducted at Livermore, Los Alamos, and Sandia National Labs. There was significant need for postprocessing software that could take full advantage of the latest and most advanced parallel processors and be able to handle models containing hundreds of millions, even billions, of nodes.

Misegades said, “EnSight Gold has proven it can handle models containing more than a billion elements, and we expect that number to extend to the tens of billions in the near future.” EnSight Gold, now the primary visualization tool for all the ASCI effort, already has been used by Los Alamos National Labs to visualize a data set containing 11.5 billion cells.

Multipipe Graphics Support

EnSight Gold provides support for multipipe graphics systems, such as those used in high-end visualization and VR environments. The application offers multipipe, single display graphics for faster rendering, multipipe rendering for multiple flat displays and PowerWall-type displays, and support for 3D input devices.

CEI is working with SGI to optimize EnSight Gold’s performance on SGI Onyx family systems with InfinitePerformance graphics. CEI uses in-house Onyx systems for its optimization efforts and, as of late, has provided kernels of EnSight Gold source code to SGI to



conduct its own optimization testing at its own facilities. InfinitePerformance is designed exclusively for interactive visualization, making it ideal for MCAE applications. InfinitePerformance allows the output of up to 16 graphics pipes to be composited into a single image for unparalleled interactive performance [up to 283 million triangles per second], providing full interactivity with even the largest and most complex MCAE simulations.

Immersive Visualization

EnSight enables users within VR environments, from desktop stereo to fully immersive rooms, to interactively interrogate the data using the application’s interactive tools. Because traditional mouse and keyboard interaction detracts from the immersive nature of VR, EnSight Gold enables 3D input [styluses, wands, cubes, and gloves] and heads-up macro [HUM], which is used for changing the viewing scene or its attributes. Typical scene changes include model part colors or visibility, starting an animation, or loading a second case.

EnSight Gold also supports multifrustum viewing, which can support virtually any multipanel display and can be used in conjunction with 3D input devices, the HUM, and parallel rendering. A view frustum defines what the application sees and displays. Desktop displays and flat screens require only a single frustum. A six-sided display, such as a Cave Virtual Automatic Environment [CAVE], requires six view frustums. Using a configuration file, EnSight Gold creates the necessary frustums, rendering only the appropriate section of the model per frustum per display surface.

EnVideo and EnLiten

More recent CEI software products include EnVideo™ and EnLiten™. EnVideo simplifies viewing videos of high-end visualizations on desktops or in virtual reality environments. EnLiten is a 3D geometry player for viewing and manipulating high-end CAE simula-

[continued on page 9]

Eye on Innovation

PARTNER

[continued from page 8]

tions. EnVideo and EnLiten files can be viewed in 3D stereo and displayed on a range of computing systems, from laptops to VR.

Automotive

EnSight applications are used in virtually every aspect of vehicle design and development by the world's major automotive brands. Nearly every Formula One racing team relies on CEI's products to aid in racecar design advancements. A number of National Association for Stock Car Auto Racing [NASCAR] teams also are recent EnSight users. Dodge Motorsports was among the first to take full advantage of EnSight postprocessing analyses through parent company DaimlerChrysler.

The global automotive engineering community also is an ardent user of Visual Area Networking [VAN], an area in which EnSight shines. In the case of U.S.-based MSXI International [MSXI], a leading supplier of engineering and technical services for automotive and other engineering-intensive industries, VAN, with the SGI® OpenGL Vizserver™ computing solution, allows the MSXI network users to receive at their desktops graphics rendered by a remote SGI® Onyx® 3000 series system with InfiniteReality3™ graphics acting as a visualization server. With OpenGL Vizserver, a team of MSXI engineers can share the advanced compute and visualization resources at their respective desktops, allowing them to perform MCAE analyses and improve collaboration and decision making.

According to Misegades, "Our experience with OpenGL Vizserver has been impressive, and our software is well matched to the concept of remote visualization. SGI technology and VAN permit easier access to EnSight's powerful visualization feature set. As most of our customers are large scientific and engineering organizations with offices in different geographic locations, remote visualization is an important part of our product offering."

Aerospace

Scientists at NASA Glenn Research Center in Ohio are designing a next-generation swirler, the device in which fuel and air are mixed within a jet engine. Its design can impact an engine's fuel efficiency, reliability, durability, and emissions, a key target of the research. To test the new design, NASA researchers are using a software code called the National Combustion Code [NCC], Pro/ENGINEER® from PTC, CFD-GEOM from CFD Research, EnSight Gold, and Silicon Graphics Onyx2 systems.

A solid geometry model of the swirler is first created in Pro/ENGINEER and then translated into initial



graphics exchange specification [IGES] format for use as a nonuniform rational B-splines [NURBS] model in CFD-GEOM, to generate the computational grid that is used to perform NCC simulations. The resulting models are as large as 10 million cells.

Four Onyx2 systems with InfiniteReality3 graphics are used in parallel mode with EnSight Gold to extract the results for visualization. In this way NASA researchers are able to perform simulation and visualization testing 5 to 10 times faster than in the past. The results are displayed on the facility's Reconfigurable Advanced Visualization Environment [RAVE™], from Fakespace.

Jay Horowitz, visualization manager at the NASA Glenn Research Center Virtual Reality Lab points out, "With EnSight Gold on the RAVE, we can now bring our data visualizations more 'out of the screen' than before. It is surprising how much better you can understand complex phenomena when you have more control over how you view them."

Future Insights

Moving forward, CEI's Misegades said work is progressing on next-generation versions of all CEI applications. Future generations of application tools also will be more modular in nature, now that the company has amassed a sizeable library of product features, or CEI visualization framework [CVF], from which it can select to create future postprocessing applications.

Developments in HPC and technology continue to rapidly advance MCAE simulation capabilities of engineers worldwide, substantially increasing the amount of data requiring postprocessing. To meet the increasing engineering needs of manufacturers, SGI will continue to work toward increasing the performance and scalability of MCAE visualization applications such as EnSight and EnSight Gold, enabling them to take full advantage of SGI HPC and visualization technologies.

Collectively, SGI and CEI will continue to innovate faster and more scalable hardware and software solutions with high-end visualization systems such as InfinitePerformance as well as future-generation, lower-cost virtual reality systems and visual workstations.



The Power To See Beyond

SGI Supercomputing and Visualization Technologies Help Advance Design, Engineering, and Manufacturing within Defense Industry

Whether at sea, by air, or on land, SGI high-performance computing (HPC) and advanced visualization technologies support the most demanding, mission-critical applications across the defense domain.

SGI HPC and visualization technologies support the defense interests of governments and their associated systems integrators by providing innovative solutions to complex product development issues while simultaneously speeding time to insight, shortening development cycles, driving down costs, and reducing time to deployment. Such capabilities provide the ability to clearly and accurately see beyond current learning and provide the power to continually reshape future product offerings.

Sea Power

Northrop Grumman Newport News, a sector of Northrop Grumman located in Newport News, Virginia, is the only company that designs, builds, refuels, and overhauls nuclear aircraft carriers, and it is one of only two to do the same for nuclear submarines. SGI workstations and high-performance computing and visualization systems have long been instrumental in helping Newport News lower the acquisition costs of its ships and reduce the total cost of ownership throughout the life cycle.

Newport News' compute and visualization power includes numerous SGI workstations and an SGI Origin 3000 system with 24 processors, eight of which are dedicated to structural analysis and 16 to 3D modeling and simulation. The system's high-performance shared memory enables Newport News to perform complex computational tasks. By including SGI® InfiniteReality® series graphics, the system also enables advanced visualization. The ability to perform these functions concurrently is unique among supercomputers and provides Newport News with a single-platform price/performance solution that can grow with the shipbuilder's ever increasing needs.

Structural Analysis

The integrity of the structure of an aircraft carrier is infinitely essential to its survival. Newport News uses SGI high-performance computing to conduct rigorous finite element analysis (FEA) on every component of the carrier structure. The analyses [static, modal, and transient] test how the ship's structure will react to every conceivable stressor once deployed, including load conditions such as hurricane-force winds, vibra-

tions from routine maneuvers, explosions above and below the waterline, and a range of other war-type events. Computer simulations under 200,000 degrees of freedom are performed on Silicon Graphics® Octane2™ workstations. The FEA-dedicated portion of the SGI Origin 3000 system handles large-scale simulation. The results of each simulation are given to the product modeler who, in turn, uses the new information to update the 3D model of the entire ship. Typical structural analysis applications used by Newport News include MSC.NASTRAN™ from MSC.Software, ABAQUS from ABAQUS, and LS-DYNA from LSTC.

In 1998, Newport News committed fully to integrating visualization technology into its shipbuilding business model by installing an SGI® Reality Center™ facility for digital prototyping, visual mock-ups, and presentations to their customers. According to Bill Kunz, Visualization Engineering Solutions project lead at Northrop Grumman Newport News, "We felt that in order to remain the world's most advanced shipyard we needed to up our competitive edge. To achieve our goal, and based on our previous successes using SGI products, we selected the SGI Reality Center facility. SGI was the only company that met all of our functional requirements."

Air Power

On October 26, 2001, the Pentagon awarded the Joint Strike Fighter (JSF) contract to Lockheed Martin. The win signaled a go-ahead for the company to produce an initial 22 F-35 aircraft as part of a U.S. Department of Defense System Development and Demonstration (SDD) phase. The \$200 billion JSF program will eventually create a fleet of some 3,000 aircraft built around three distinct design specifications—conventional takeoff and landing, carrier variant, and short takeoff/vertical landing.

All three JSF designs will share key high-cost components, including propulsion systems, and have a "cost commonality" of 70% to 90%. The goal is to realize significant cost savings in the development, manufacture, and maintenance of the aircraft.

Tom Burbage, executive vice president and general manager, Lockheed Martin JSF Program, said, "Our JSF approach [enabled by SGI visualization technology] will radically reduce the cost of sustaining U.S. airpower by ensuring affordability during SDD, production, operations, and support and by achieving operational excellence throughout the program."

[continued on page 11]



"We felt that in order to remain the world's most advanced shipyard we needed to up our competitive edge. To achieve our goal, and based on our previous successes using SGI products, we selected the SGI Reality Center facility. SGI was the only company that met all of our functional requirements."

— Bill Kunz
Visualization
Engineering Solutions
Project Lead, Northrop
Grumman Newport News

Eye on Innovation

INDUSTRY



[continued from page 10]

Recently, Lockheed Martin Aeronautics Company added more SGI horsepower to its flight simulation lab in Fort Worth, Texas—specifically, SGI Onyx 3000 series and SGI® Onyx® 300 visualization systems, as well as Silicon Graphics Octane2 workstations. Lockheed Martin engineers utilize the compute and graphics power of SGI systems to tweak the internal design of the F-35. The Fort Worth facility, along with a second flight-simulation facility in Palmdale, California, supports design and evaluation of the flight controls and aerodynamics, as well as the integrated tactical performance of JSF aircraft.

“The cornerstone of the F-35 is affordability, achieved in large part through a very high level of common parts and systems across the three versions of the aircraft,” said Burbage. “SGI visualization technology has helped us achieve this commonality through significant advancement in the design of the aircraft, and we will rely on that technology as we approach the JSF program’s next milestone, the Preliminary Design Review, scheduled for March 2003.”

To date, a number of F-35 external design changes have been finalized with the help of SGI visualization technology. Among them was raising the top surface of the aircraft slightly along the centerline, a change that helped to increase fuel capacity by 300 pounds, thereby extending the range of the JSF.

Propulsion Optimization

To power the JSF aircraft developed by Lockheed Martin, the Pentagon selected the Pratt & Whitney F135 engine, a derivative of the company’s F119 engine now used in the military’s F-22 fighter aircraft.

Pratt & Whitney Canada, the world’s leading manufacturer of gas turbine aircraft engines, utilizes SGI HPC for structural FEA and computational fluid dynamics [CFD] modeling to ensure accuracy and performance in the engine design. Using SGI HPC and solver technology, the company has developed full-size models of JSF engines for CFD evaluation. Using the models will reduce engine development costs, allow for the direct transfer of engine design information to manufacturing, and significantly reduce performance certification time.

Ground Combat Power

The mission of the U.S. Army Research Laboratory [ARL] is to provide innovative science, technology, and analyses to enable full-spectrum operations. To further its ability to study how soldiers use equipment in combat zones, the ARL recently installed a RAVE™ II [Reconfigurable Advanced Visualization Environment II] stereographic display system from Fakespace

Systems at its Tactical Environment Simulation [TES] Facility located in Aberdeen Proving Ground, Maryland.

Within the TES Facility, an Immersive Environment Simulator integrates the RAVE II system with motion tracking and an omnidirectional treadmill to allow soldiers to literally run and move in any direction within virtual hostile terrain and combat conditions. It is the first system of its type to include an omnidirectional treadmill, and it also has an extremely accurate inertial-acoustic position measuring system, which enables visualizations to track in real time with the user’s changing point of view. Powering the Immersive Environment Simulator is a three-pipe, 16-processor Onyx 3000 series system with an InfiniteReality3 graphics subsystem. The ARL selected the Onyx system, in part, for its ability to accommodate, simultaneously, the RAVE II system’s three self-contained 10x12.5-foot rear-projected modules and the 44-speaker, 155-decibel sound system that is part of the Immersive Environment Simulator.

The ARL sees the new TES facility as an exceptional resource for conducting human factors research on pre-prototype hardware and software, such as new helmet designs and target acquisition systems. More important, by applying the most advanced technology to the study of soldier performance under highly controlled conditions, the U.S. Army will enable discoveries that increase the safety and effectiveness of its battlefield forces.

Forward Progress

To ensure that defense programs continue to realize their computational and visualization technology benchmarks, SGI will continue to develop high-performance computing and advanced visualization solutions that continue to expand the boundaries of design, development, and manufacture for a diversity of vital programs.



MCAE Applications

[continued from page 6]

The capabilities of SGI Altix 3000 create opportunities for increased innovation, such as expanded use of time-dependent and time-accurate MCAE simulations. Additionally, there are opportunities for more production-level coupling of applications for multidisciplinary design optimization. The shared-memory, multi-purpose architecture of SGI Altix 3000 is a platform that is well suited to the simultaneous demands of all three MCAE disciplines. SGI Altix 3000 provides the ability to capture more realism in simulations.

The record-setting performance levels of the SGI Altix 3000 systems will be at the forefront of performance

because SGI designed the Altix 3000 servers with the ability to upgrade to the future Itanium 2 architectures code-named Madison and Montecito. Perhaps of even greater significance is the fact that SGI has accomplished these technological breakthroughs by working closely with the Linux open-source community and continues to offer much of its work back to the community, to be incorporated into later versions of the Linux kernel. These investments will ensure success of the new SGI Altix 3000 technology and the company's continued commitment to the MCAE community in delivering valuable leadership in HPC advancements for the manufacturing industry.

Procter & Gamble

[continued from page 4]

tem, P&G first established performance benchmarks using a system that belonged to SGI. We then had SGI confirm the benchmarks prior to our decision to purchase. After SGI finished the installation, we again had them confirm the benchmarks on our new system. One of the most important elements of P&G is reliability engineering. It's a discipline we understand very well."

Collaboration

Within P&G's technical community, collaborations between product designers and the people who build the manufacturing systems are well established. However, one of the interesting things that high-performance computing has provided P&G is the development of a more common language among a smaller subgroup of the technical community—those who normally are responsible for determining the functionality of a product and the equipment and technology that make those products.

"High-performance computing has ended what used to be the distrust of the esoteric," says Lange. "Now when

we say, 'We ran your bottle on a virtual test and here's what happened to it, how do you think you can fix that?' designers are suddenly interested. They see more clearly that they can preserve the intent of the package design and still meet the necessary functional requirements of the product. In fact, what I've seen in the packaging community is almost an excitement to help the design community create more exciting designs and, at the same time, preserve the economics of the package."

Summing it Up

P&G has documented proof that a number of product successes have directly resulted from using CAE tools and supercomputers to create them. The use of CAE and supercomputers like the SGI Origin 3800 system have created an entirely new way of doing engineering at P&G. "We now possess complex problem-solving capabilities that were unthinkable 10 years ago," says Lange. "These tools allow us to be more innovative, more creative, and enable us to develop products in less time and at lower costs than ever before. But the real promise they hold for P&G and the package goods industry is the ability to create better products for consumers."

Don't miss the exclusive SGI Eye on Innovation Webcasts for Manufacturers, hosted by Daratech, Inc. For dates, times, and more information visit: <http://sgievent.com/eyeoninnovation>



Corporate Office
1600 Amphitheatre Pkwy.
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America | (800) 800-7441
Latin America | (52) 5267-1387
Europe | (44) 118.925.75.00
Japan | (81) 3.5488.1811
Asia Pacific | (65) 6771.0290

If you wish to subscribe to this newsletter via email, please visit www.sgi.com/subscribe

©2003 Silicon Graphics, Inc. All rights reserved. Silicon Graphics, SGI, Origin, Onyx2, Onyx, IRIX, OpenGL, InfiniteReality, Octane, and the SGI logo are registered trademarks and NUMalink, SGI ProPack, InfinitePerformance, OpenGL Vizserver, InfiniteReality3, Octane2, and Reality Center are trademarks of Silicon Graphics, Inc., in the United States and/or other countries worldwide. MIPS is a registered trademark of MIPS Technologies, Inc., used under license by Silicon Graphics, Inc., in the United States and/or other countries worldwide. Intel and Itanium are registered trademarks of Intel Corporation. Linux is a registered trademark of Linus Torvalds. UNIX is a registered trademark of The Open Group in the U.S. and other countries. All other trademarks mentioned herein are the property of their respective owners. Images courtesy of Procter & Gamble, Computational Engineering International, NASA, Northrop Grumman Newport News, U.S. Army Research Laboratory.